

UNCLASSIFIED

Defense Technical Information Center
Compilation Part Notice

ADP023796

TITLE: HPC Access Using "KVM over IP"

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Proceedings of the HPCMP Users Group Conference 2007. High Performance Computing Modernization Program: A Bridge to Future Defense held 18-21 June 2007 in Pittsburgh, Pennsylvania

To order the complete compilation report, use: ADA488707

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP023728 thru ADP023803

UNCLASSIFIED

HPC Access Using “KVM over IP”

Dan Kedziorek and Gregory P. Czerniak

*US Army Research, Development and Engineering Command, Tank-Automotive Research,
Development and Engineering Center (RDECOM-TARDEC), Warren, MI
{dan.kedziorek, gregory.p.czerniak}@us.army.mil*

Abstract

The Research, Development and Engineering Command, Tank-Automotive Research, Development and Engineering Center (RDECOM-TARDEC) has been working with IP Video Systems (formerly Teraburst Networks) since 2002. Among the various keyboard-video-mouse over internet protocol (KVM over IP) solutions that are on the market today, our research has revealed that the IP Video Systems solutions seem to be the best suited for Department of Defense (DoD) needs, especially those of the High Performance Computing Modernization Program (HPCMP). Recognizing this, in late 2003, we submitted an Army Small Business Innovative Research (SBIR) topic to tailor further development of KVM-over-IP technology to the needs of high performance computing centers. These requirements included low-latency transmission of high resolution and high refresh rate video, active stereo synchronization, peripheral, and audio signals. Security, advanced compression techniques, configurability, and a reduced form factor were other considerations. Expected completion of the SBIR Phase II effort is planned for late FY 07. This paper will discuss the history, success, technical specification, and future plans of this technology as a result of this SBIR effort.

1. Introduction

A persistent challenge in the high performance computing (HPC) community is how to provide remote visualization capability to its users. The RDECOM-TARDEC HPC Group has addressed this problem locally within its installation through a legacy system—an elaborate network of Lightwave VDE/200 KVM-over-Fiber (Keyboard, Video, and Mouse) devices installed throughout the TARDEC campus. Implementation of this system required the deployment of dedicated multi-mode fiber optic lines to serve as the transmission medium. This solution, while effective within the boundaries of TARDEC, fails to address the issue of long-distance

transmissions of KVM data, as multi-mode fiber can only span about two miles before signal degradation occurs. Although single-mode fiber KVM-over-Fiber devices exist that are capable of longer distance transmissions, the high cost of these devices coupled with the difficulties—and cost—of installing dedicated long-distance fiber optic lines between installations make this an impractical solution for widespread deployment. Aggravating the situation further is the fact that dedicated fiber is best suited for installation for point-to-point communications.

A more dynamic and economical solution is a KVM-over-IP technology, which uses a pre-existing Transmission Control Protocol/Internet Protocol (TCP/IP) network to transmit KVM data between two locations. However, the level of performance and functionality present in the current consumer KVM-over-IP devices makes them less than desirable for Department of Defense (DoD) HPC applications.

To address the specific needs of the DoD HPC community, the TARDEC HPC Group undertook a three-year development effort through the pursuit of an Army-funded Phase-II Small Business Innovative Research (SBIR) effort with IP Video Systems (formerly known as TeraBurst) to produce an enhanced version of their V2D product (henceforth referred to as “V2D-SBIR”). At the time of this paper’s publication, the SBIR will be near completion. Additional testing and future demonstrations are expected subsequent to this paper.

2. Functionality

To accommodate remote use of the high-end visualization capabilities of a DoD HPC facility, many advanced features are necessary. TARDEC’s -HPC SBIR with IP Video Systems indicated specific requirements for creating a KVM-over-IP device that could be used for HPC visualization purposes. These requirements included support for a universal serial bus (USB) keyboard and mouse, multi-channel digital audio, full-duplex Recommended Standard 232 (RS232) transmissions, and receiver-side graphic genlock support.

A significant requirement for a modern KVM-over-IP device is USB keyboard and mouse support over the USB Human Interface Device standard. At the time of the SBIR's initial signing, few consumer-level KVM-over-IP products supported the use of USB devices, instead preferring the aging PS/2 interface. To address this, adding a USB keyboard and mouse control was a specific requirement in the SBIR contract, and TARDEC engineers have verified that this functionality now exists within the V2D-SBIR. Incorporating a generalized USB subsystem that would support the transmission of any USB signal was beyond the scope of this development effort. Therefore, the SBIR-enhanced V2D is capable of transmitting USB keyboard and mouse signals over IP, but it does not support USB motion trackers or other USB devices at this time. This capability may be addressed in a later version of the V2D-SBIR system.

Another requirement for widespread KVM-over-IP use within HPC facilities is support for multi-channel digital audio. This is especially necessary for supporting visualization systems such as Cave Automatic Virtual Environment (CAVE)-like environments that utilize three-dimensional (3D) spatial audio. High performance computing visualization environments typically use Alesis Digital Audio Tape (ADAT) audio signals for transmitting digital audio from supercomputers. Therefore, ADAT audio support with a minimum of four channels was included as a requirement for the SBIR effort. TARDEC engineers have verified that the device can transmit at least one channel of audio over the ADAT port, and a full four-channel test will be conducted in the near future.

Visualization environments such as CAVE-like displays use a variety of motion tracking devices for user interaction. The majority of these devices use an RS232 serial interface, making it a necessity for virtual reality and visualization capabilities. To create a KVM-over-IP system tailored for HPC visualization tasks, TARDEC included RS232 support as a requirement in the SBIR effort. As a result, the V2D is now capable of full-duplex transmission of serial data over the RS232 port. The functionality has been verified by IP Video Systems by performing a serial file transfer between two computers using the serial subsystem of the V2D-SBIR.

Unlike most users of a KVM-over-IP system, a major need for DoD HPC users is the ability to synchronize the output of multiple video displays. Since practically all CAVE-like displays use multiple video outputs to produce a composite image, a means to synchronize the screen refresh rates through support of genlock is necessary to produce a correct immersive 3D environment using active stereo. IP Video Systems was successful in incorporating this capability into the V2D-SBIR device. Through internal tests, they have successfully synchronized six monitors with the V2D developed under

this SBIR. TARDEC experiments with the V2D-SBIR's genlock system will occur once TARDEC receives all of the deliverables upon completion of the SBIR.

3. Performance

In order to transmit video over IP at an acceptable performance level, the V2D-SBIR applies lossy compression to the signal. To achieve this, the V2D-SBIR device features a field programmable gate array (FPGA) chip to implement their proprietary compression algorithms in hardware, which allows for very high-speed compression of video data. Lossy compression implies degradation in image quality, but due to the limited bandwidth of most conventional TCP/IP networks, a compromise must be made between image quality and frame rate. Given the subjective nature of image quality, this compromise can only be made by the end user. Despite the burden for the user to determine the compression and speed balance needed for their specific application, most KVM-over-IP devices on the market do not provide options for configuring the amount of compression being applied to the data stream. In contrast, the V2D-SBIR offers the user an extensive level of configuration options for finding a perfect balance.

The V2D-SBIR hardware differentiates between lossy compression performed on parts of the display image that stay mostly static (such as large parts of a typical desktop in a graphical user interface) and on parts of the display image that change rapidly (such as full-motion video). The amount of lossy image compression performed on the mostly static areas of the screen is determined by a user-defined coefficient labeled "Low Compression", which is represented as an integer between zero and ten. Likewise, the amount of image compression performed on the more dynamic areas of a display image is labeled "High Compression", and its intensity is also represented by an integer coefficient between zero and ten.




TARDEC HPC engineers have developed several in-house automated benchmarking tools that allow the team to measure the performance of the V2D-SBIR hardware. One of these tools systematically sets the hardware to all 121 combinations of Low and High Compression for five seconds each and measures the average bandwidth usage for that duration. A full table of these bandwidth values can be created in just over ten minutes, during which a consistent source of animation is continuously sent over the device. The frame rate of the animation is ideally locked at its maximum value (the refresh rate of the video signal). If there is insufficient bandwidth, the remote display frame rate will suffer, resulting in undesirable graphic artifacting.

In obtaining the data contained in Tables 1 and 2, the V2D-SBIR hardware had 100 megabits of bandwidth available over the internal TARDEC network, which is the maximum speed of the V2D's network card. Therefore, if the device is using less bandwidth than the maximum (the maximum was considered to be 85 Mb/sec in this test due to transmission overhead), it is assumed that the video is displaying at full speed.

For this paper, two tests with two different video inputs were performed. One video input was a five-second excerpt from the 720p BBC Motion Gallery's "Andes" video available on Apple.com running in a continuous loop (to synchronize with the bandwidth sampling period), which was scaled and letterboxed to 1024x768 at 60Hz. This was chosen to be a realistic demonstration of real-time video streaming for the V2D-SBIR hardware. The other video input used was the Windows XP "Beziers" screensaver at 1024x768 at 60Hz with all settings at maximum, which was determined to be extremely demanding of the V2D's compression system. This was chosen to be a demonstration of a worst-case yet realistic scenario that the V2D may have to encounter in regular use.

Tables 1 and 2 provide a summary of the amount of bandwidth necessary to send two different types of video, each with a different level of general information entropy, at every combination of compression settings offered by the V2D-SBIR hardware, assuming a constant frame rate. For both of these tests, the benchmarking program was executed three times and the results were averaged into one table. Note that there are, in fact, settings available that will allow for a very usable interactive session at fairly low bandwidths. The following is the color code legend for the tables. (Refer to the full-color tables at the end of this paper.)

Color-Coding Legend

	Can be sent over a DS3 (<45 Mb/sec)
	Can be sent over LAN (Between 45 Mb/sec and 80 Mb/sec)
	Cannot be sent (>80 Mb/sec)

4. Networking

The goal of the V2D-SBIR project is to transfer data from a transmitter to a receiver over any IP-based network. However, due to security and performance practicalities regarding conventional networks (firewalls, Network Address Translations, Quality of Service issues, etc), one must also take into account the network that the V2D-SBIR will be transmitting data through.

As an example of these complications, this paper will cover the basic network issues encountered during

TARDEC-HPC's first Internet test of the V2D-SBIR which was conducted between the Detroit Arsenal in Warren, MI, and its partner, Automation Alley, headquartered in Troy, MI (approximately nine miles away). As TARDEC and Automation Alley are two separate organizations, their networks are independent, connected to each other only via the Internet. The relatively-short distance between installations and the network separation made the Automation Alley headquarters an excellent choice for our first Internet test of the V2D-SBIR. Figure 1 depicts a logical network diagram used for these tests.

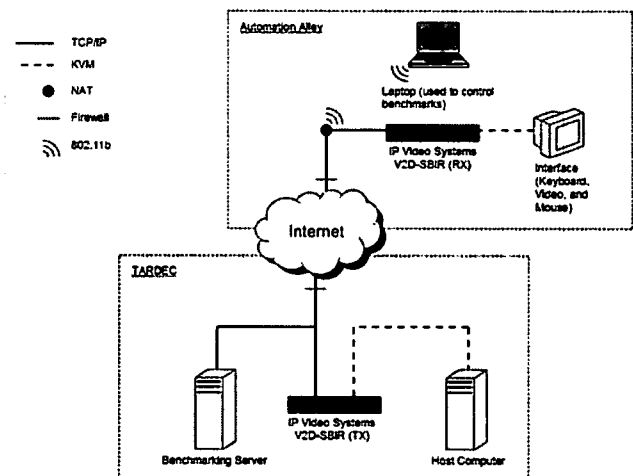


Figure 1. Logical network diagram for the TARDEC/Automation Alley tests

One major complication in the use of a KVM-over-IP device is the firewall protection between Internet networks. The TARDEC/Automation Alley test had many coordination and configuration issues with firewalls before the system could communicate reliably. On both networks, network administrators had to forward TCP and Universal Datagram Protocol (UDP) traffic for certain ports. Making the problem even worse were port blocks imposed by network authorities above the TARDEC command level, which caused traffic to be blocked even when the ports were supposedly forwarded locally. After realizing the higher-level issue, TARDEC HPC successfully found an authorized port that worked between the two networks, which allowed for data transmission between the two facilities.

Another complication in the TARDEC/Automation Alley test was bandwidth limitations. The Automation Alley network provided an Internet connection that was able to communicate consistently at only 2.5 megabits per second. This led to significantly reduced frame rates, especially at a resolution of 1,024x768. However, the speed of the transmission is still fast enough to support typical desktop usage, although full motion video at 2.5

Mbps requires extremely high compression levels to obtain a usable frame rate of over 10 frames per second.

When considering the use of KVM over IP technologies, another networking issue that must be addressed is that of latency and packet loss. During the TARDEC/Automation Alley test, there was relatively little network latency, and packet loss was only an occasional problem. However, the hardware is capable of working with even the one-second latency, but would make any interactive use of the V2D effectively impossible. Therefore, obtaining favorable latency and low packet loss between the target networks is a prerequisite for proper V2D-SBIR operation.

In terms of security, the V2D-SBIR offers secure shell access to its configuration menu system for secure login, as well as a serial console and a direct monitor/keyboard connection. The V2D-SBIR does not offer on-board encryption beyond its compression algorithms for the video feed or the keyboard/mouse/audio information. In order to facilitate a secure transmission of V2D-SBIR data over the Internet, IP Video Systems recommends the use of dedicated hardware VPN devices that can transmit over UDP.

Although KVM-over-IP devices are connected via IP networks, problems encountered during the Automation Alley test were relatively minor. Though when facing a larger deployment of V2D-SBIR hardware, especially for a site that may have a centralized bank of transmitters, the coordination of port settings and IP address assignment would become essential.

5. Conclusion

Overall, the V2D-SBIR project has led to great advances in KVM-over-IP technology, particularly in the realm of DoD High Performance Computing expectations. However, just as technology continues to progress, there will always be room for improvement in the KVM-over-IP capabilities.

One improvement TARDEC HPC would like to see is a generalized USB transmission system capable of transmitting any USB data over the network rather than just the USB keyboard-and-mouse-only configuration that

is on the current V2D-SBIR prototype. Due to the extremely high bandwidth requirements of many USB devices, it may be impractical to create such a system with the current levels of bandwidth available.

Another desired improvement is Gigabit Ethernet support. Although only the expensive, high-bandwidth Internet connections are currently capable of transmitting more than 100 megabits per second of data, preparing for the future of extremely high-bandwidth Internet connections is a wise decision. As KVM-over-IP devices are bandwidth-intensive devices, this progression would allow for much higher resolution video at steadier frame rates over the proper connections.

Another possibility for additional research is the concept of a small form-factor version of a KVM-over-IP device, specifically one that could run within the PCI bus of a workstation. Such a system would allow a wide variety of computers to feature full KVM-over-IP support without the installation of rack-mount V2D-SBIR devices. This step would require additional funding and further research, as the small form-factor version would be such a radical change over the rack-mount system that a large amount of architecture redesign would be necessary.

Regardless of the desired improvements, the IP Video Systems V2D-SBIR prototype exposes DoD HPC engineers to a new horizon of KVM-over-IP transmission of visualization and virtual reality environment data. As the cost of Internet bandwidth goes down and the cost of dedicated fiber goes up, the KVM-over-IP solution, with its more dynamic connection methodology, will have a definite place within the DoD HPC community.

The TARDEC HPC Group and IP Video Systems are looking for potential DoD HPC users that could benefit from this remote visualization capability. We can provide the hardware and apply our testing experience to demonstrate that the capability does exist and meets the needs of the DoD HPC community. We also are looking for opportunities and customers that would support us in continuing the development of this capability by means of a SBIR Phase III. Contact the TARDEC HPC Helpdesk for further details at hpc@tacom.army.mil.

Table 1. BBC motion gallery "Andes" excerpt, 1,024x768@60Hz

Bandwidth used in Mb/sec (assuming maximum framerate)											
Low\High	0	1	2	3	4	5	6	7	8	9	10
0	75.6	45.4	44.7	18.6	18.0	17.9	13.3	12.9	9.5	8.4	7.8
1	75.3	44.9	43.6	17.6	16.7	16.1	11.9	11.2	7.7	6.8	6.2
2	74.8	44.1	44.0	17.4	16.4	16.0	11.6	10.9	7.4	6.3	5.8
3	75.5	43.8	42.8	16.3	15.6	14.9	10.7	10.2	6.7	5.6	5.1
4	75.5	43.3	42.9	16.3	15.5	15.0	10.6	10.0	6.6	5.5	5.1
5	74.8	43.2	43.0	16.3	15.5	14.8	10.6	10.1	6.6	5.5	5.1
6	75.7	43.3	42.6	16.3	15.5	14.9	10.5	10.0	6.6	5.4	5.0
7	75.1	43.4	42.5	16.3	15.4	14.7	10.4	10.1	6.6	5.4	5.0
8	74.7	43.2	42.0	16.3	15.5	14.9	10.4	10.0	6.5	5.3	4.9
9	75.8	43.2	42.6	16.2	15.4	14.8	10.2	10.0	6.5	5.3	4.9
10	74.0	43.0	42.0	16.1	15.5	14.9	10.3	9.9	6.4	5.3	4.8

Table 2: Windows XP "Beziers" screensaver, 1,024x768@60Hz.

Bandwidth used in Mb/sec (assuming maximum framerate)											
Low\High	0	1	2	3	4	5	6	7	8	9	10
0	90.7	90.4	90.4	90.4	89.9	89.6	88.5	88.9	86.6	79.3	78.2
1	90.7	90.6	90.3	89.5	90.2	89.9	89.2	88.4	84.5	73.1	69.6
2	90.9	90.6	91.0	90.7	90.4	90.7	89.8	89.4	86.8	76.3	68.5
3	90.7	90.1	90.7	88.8	89.1	89.6	82.7	78.4	52.9	42.7	43.7
4	90.4	90.0	89.8	87.0	84.8	86.2	84.0	81.4	69.6	52.2	45.6
5	90.9	90.6	90.7	89.3	89.3	89.2	87.2	87.5	78.8	65.6	55.2
6	90.3	91.0	90.7	89.3	87.7	88.7	85.6	84.6	75.1	52.1	36.3
7	90.5	90.1	90.2	87.1	86.5	84.6	65.9	60.9	42.8	30.2	31.2
8	90.3	88.8	87.1	80.3	80.6	81.0	75.9	74.3	42.1	22.3	17.4
9	90.3	90.2	89.9	87.8	83.8	86.3	77.8	78.2	58.2	40.5	23.0
10	90.3	90.2	89.9	83.2	85.9	86.2	78.4	77.6	42.6	19.5	10.6